APPLICATION FOR UNITED STATES LETTERS PATENT

INVENTORS: Maged M. GHONEIMA, Peter W. CAPUTA, Muhammad M.
               KHELLAH, Ram KRISHNAMURTHY, James W. TSCHANZ,
               Yibin YE, Vivek K. DE, and Yehea I. ISMAIL

TITLE: SINGLE- STAGE AND MULTI-STAGE LOW POWER
               INTERCONNECT ARCHITECTURES

SINGLE-STAGE AND MULTI-STAGE LOW POWER INTERCONNECT ARCHITECTURES

5 ## FIELD

Embodiments of the present invention relate to interconnects. More particularly, embodiments of the present invention relate to reducing energy consumption of single-stage and/or multi-stage interconnect architectures.

10 ## BACKGROUND

Low power design is a major issue in integrated circuit design. With the continuous scaling of technology, increased die area and faster clock speeds, the power consumption of on-chip interconnects is very significant.

Long on-chip interconnects are one of the main bottlenecks in microprocessor

15 design due to the impact of coupling capacitance on both interconnect (or bus) delay and power dissipation. For example, in deep submicron technology, the aspect ratio may be over 2.0 for intermediate wiring layers. This indicates that the lateral component of the interconnect capacitance (i.e., coupling capacitance) continually grow to dominate the total interconnect capacitance due to reduction in wire pitch

20 and an increase in the aspect ratio of the interconnect. Lateral interconnect capacitance components are from three to five times as much as vertical capacitance components. Architectures should be designed to reduce the impact of the coupling capacitance on energy dissipation and delay of on-chip interconnects.

Introducing a relative delay between adjacent lines reduces the maximum

25 effective coupling capacitance by approximately half which leads to a considerable

reduction in the maximum total capacitance and hence bus delay. Techniques may introduce a relative delay between oppositely switching lines to reduce interconnect delays and increase bus performance. A delay may be introduced by adding extra buffers on alternate bus lines, or skewing the inverters to change the rise and fall

5   delays of the repeaters. On average, if a relative delay is permanently introduced between any two adjacent lines, then the average energy dissipation of the bus will not change.

BRIEF DESCRIPTION OF THE DRAWINGS

10   The foregoing and a better understanding of the present invention will become apparent from the following detailed description of arrangements and example embodiments (and the claims) when read in connection with the accompanying drawings, all forming a part of the disclosure of this invention. While the foregoing and following written and illustrated disclosure focuses on disclosing

15   arrangements and example embodiments of the invention, it should be clearly understood that the same is by way of illustration and example only and the invention is not limited thereto.

The following represents brief descriptions of the drawings in which like reference numerals represent like elements and wherein:

20   FIG. 1 is a block diagram of a computer system according to an example arrangement;

FIG. 2A illustrates a two-line directional interconnect structure according to an example arrangement;

FIG. 2B illustrates the delay with three types of simultaneous transitions on the two-line directional interconnect structure of FIG. 2A according to an example arrangement;

FIG. 3 shows drivers for two adjacent signal lines according to an example

5   embodiment of the present invention; and

FIG. 4 is a multi-stage interconnect architecture according to an example embodiment of the present invention.

DETAILED DESCRIPTION

10   In the following detailed description, like reference numerals and characters may be used to designate identical, corresponding or similar components in differing figure drawings. Further, in the detailed description to follow, example sizes/models/values/ranges may be given although the present invention is not limited to the same. Where specific details are set forth in order to describe example

15   embodiments of the invention, it should be apparent to one skilled in the art that the invention can be practiced without these specific details.

In this disclosure, the terminologies interconnect and bus may be used interchangeable to relate to a plurality of signal lines in which each signal line may separately carry separate signals. The terminology line, signal line or interconnect

20   may relate to one of the signal lines on the interconnect or bus. Additionally, adjacent signal lines relate to signal lines that are provided immediately adjacent to one another. For ease of illustration, the examples and the figures discussed herein may only describe and/or show a small number of signal lines on an architecture.

Other numbers of signal lines are also fully within the scope of this disclosure. Additionally, while the description may relate to signals traveling in one direction, the interconnects may be bi-directional.

Interconnect architectures will hereafter be discussed with respect to reducing

5   the power consumption of single stage and multi-stage buses, with minimal delay penalties.

Embodiments of the present invention may introduce a relative delay between adjacent interconnects to not only reduce a worst-case delay, but also reduce an average interconnect energy dissipation. This is performed by reducing the drivers'

10 voltage supply of alternate bus lines, such that for any two adjacent bus lines, the driver voltage supply of one line has the nominal value V_{DD}, while the driver voltage supply of the other line has a reduced voltage (V_{DD} -ΔV). Reducing the voltage supplies of alternate bus lines increases their delay compared to their adjacent lines and thus introduces a relative delay between adjacent signal lines. Relative delays

15 between adjacent lines reduce the effective coupling capacitance which in turn further reduces the overall interconnect delay and hence improves the bus performance. Relative delays do not affect the average energy dissipation, but lowering the alternate line voltage supplies lowers their average energy dissipation and hence reduces the overall bus energy dissipation.

20   Embodiments of the present invention may combine low power bus coding followed by dynamic delaying the data of certain signal lines of a multi-stage interconnect structure. Low power bus coding encodes the transmitted data on the signal lines to lower the worst-case effective coupling capacitance and hence reduce

4

the average energy dissipation. Dynamic delaying introduces relative delays between specific lines according to their relative switching activity in order to reduce the worst case coupling capacitance and hence the overall average energy dissipation. Low power bus encoding and dynamic delaying can be cascaded as

5   their operation is independent of one other and thus their combined energy dissipation can be achieved. The impact of delay penalty associated with the bus encoding/decoding hardware and the dynamic delay circuit can be reduced by amortizing it over several stages when this architecture is applied to multi-stage time-borrowing buses.   Thus, combining low power encoding with dynamic delaying

10   in multi-cycle time-borrowing buses leads to significant energy savings while adding a small delay overhead per bus stage.

A General Computer System

FIG. 1 is a block diagram of a computer system 100 according to an example

15   arrangement. Other arrangements are also possible. A computer system 100 may include a processor 105 coupled to a processor bus 110. The processor 105 may include a first level (L1) cache memory (not shown in FIG. 1).

The processor 105 may also be coupled to a cache memory 107, which is a second level (L2) cache memory, via a dedicated cache bus 102. The L1 and L2

20   cache memories can also be integrated into a single device. Alternatively, the cache memory 107 may be coupled to the processor 105 by a shared bus.

A chip set 120 may also be coupled to the processor bus 110. The chip set 120 may include a memory controller for controlling a main memory 113. Further,

the chip set 120 may be coupled to a video device 125 that handles video data requests to access the main memory 113. The video device 125 may include a video monitor such as a cathode ray tube (CRT) or liquid crystal display (LCD) and necessary support circuitry.

5   The main memory 113 may be coupled to the processor bus 110 through the chip set 120. The main memory 113 and the cache memory 107 may store sequences of instructions that are executed by the processor 105. The main memory 113 may include a Rambus dynamic random access memory (RDRAM) system or the main memory 113 may have other configurations. The sequences of

10   instructions executed by the processor 105 may be retrieved from the main memory 113, the cache memory 107, or any other storage device. Additional devices may also be coupled to the processor bus 110, such as multiple processors and/or multiple main memory devices. The computer system 100 is described in terms of a single processor; however, multiple processors can be coupled to the processor bus

15   110.

      The processor bus 110 may be coupled to a system bus 130 by the chip set 120. The system bus 130 may be a Peripheral Component Interconnect (PCI) bus adhering to a Specification Revision 2.1 bus developed by the PCI Special Interest Group of Portland, Oreg.; however, other bus standards may also be used. Multiple

20   devices, such as an audio device 127, may be coupled to the system bus 130.

      A bus bridge 140 may couple the system bus 130 to a secondary bus 150. For example, the secondary bus 150 may be an Industry Standard Architecture (ISA) Specification Revision 1.0a bus developed by International Business Machines

of Armonk, N.Y. However, other bus standards may also be used such as Extended

Industry Standard Architecture (EISA) Specification Revision 3.12 developed by

Compaq Computer, et al. Multiple devices, such as a hard disk 153 and a disk drive

154 may be coupled to the secondary bus 150. Other devices, such as cursor

5 control devices (not shown in FIG. 1), may be coupled to the secondary bus 150.

The computer system 100 described above (or other electronic

devices/systems) may include one or more chips (also referred to as chipsets or

dies) within various components or forming one or more of the components. In

disadvantageous arrangements, long interconnects traversing across a die (or

10 between dies) incur a substantial delay due to large RC (i.e., resistance x

capacitance) values that impede a signal traveling from a source to a destination. In

deep submicron (DSM) technologies, these RC values have dramatically increased

due to the dominating effect of coupling capacitance. The growing impact of the

coupling capacitance between adjacent interconnects affect the maximum speed at

15 which an entire chip may operate, as well as the overall energy dissipation.

Repeater circuits may be provided to aid in these problems. Embodiments of the

present invention relate to on-chip interconnect architectures that reduce the

average interconnect dissipation as will now be described.

20 **Effect of Capacitive Coupling on Interconnect Delay and Energy Dissipation**

FIG. 2A illustrates a two-line directional interconnect structure according to an

example arrangement. FIG. 2B illustrates the delay with three types of simultaneous

7

transitions on the two-line directional interconnect structure of FIG. 2A according to an example arrangement. Other arrangements and delays are also possible.

More specifically, FIG. 2A shows an interconnect architecture 200 having drivers 210 and 220, interconnects 240 and 250 and receivers 270 and 280. A

5 signal I1 may be received by the driver 210 and transmitted along the interconnect 240 to the receiver 270, which outputs signal O1. Likewise, an input signal I2 may be received by the driver 220 and transmitted along the interconnect 250 to the receiver 280, which outputs signal O2. In this figure, C_g represents the capacitance between one of the signal lines and a substrate, and C_c represents the capacitance

10 between adjacent signal lines. The inter-wire coupling capacitance C_c is the more dominant capacitance in current and future deep submicron (DSM, hereafter) technologies, which severely increases interconnect delays, interference and power consumption.

The propagation delay across an interconnect is due to the finite time it takes

15 to charge or discharge its capacitive components. Thus, the presence of a coupling capacitor C_c between any two coupled lines L_1 and L_2, will increase their propagation delays due to the time required to charge/discharge the coupling capacitor C_c by $\Delta Q_c = C_c \times (\Delta V_1 - \Delta V_2)$, where ΔV_1 and ΔV_2 are the voltage level variations on lines L_1 and L_2, simultaneously. The presence of a coupling capacitor also increases the

20 energy dissipation due to the losses associated with its charging and discharging. The average energy dissipation in each switching incident due to capacitive coupling is $E_C = 0.5 \times C_c \times (\Delta V_1 - \Delta V_2)^2$.

FIG. 2B shows the switching activity of two coupled interconnects at three separate times (a), (b) and (c). The voltage levels on each interconnect switch between 0, representing a LOW logic level, and V_{DD}, representing a HIGH logic level. At time (a), switching occurs on both interconnects 240 and 250 in the same

5   direction. At time (b), the signal traversing the interconnect 250 switches alone while the other interconnect 240 remains quiet. At time (c), the signal on the interconnect 250 switches in a direction opposite to that of the signal on its immediately adjacent interconnect 240.

When only one interconnect 250 switches (i.e. time (b)), then the voltage level

10   variations of the interconnects are $\Delta V_1 = 0$ and $\Delta V_2 = -V_{DD}$. Thus, the coupling capacitor is charged by $|\Delta Q| = C_c \times |\Delta V_1 - \Delta V_2| = C_c V_{DD}$ leading to a certain propagation delay (T_b) for the switching line. The energy dissipation due to capacitive coupling in this switching scenario is $E_{cb} = 0.5 \times C_c \times (\Delta V_1 - \Delta V_2)^2 = 0.5 \times C_c \times V_{DD}^2$.

The best-case scenario for delay and energy dissipation occurs when the

15   interconnects simultaneously switch in the same direction (i.e., time (a)), and thus both have the same voltage level variations $\Delta V_1 = \Delta V_2 = V_{DD}$. The coupling capacitor in this case, does not increase the delay or energy dissipation of the switching lines, as the coupling capacitor is neither charged nor discharged $(\Delta V_1 - \Delta V_2) = 0$, and hence the total delay of the signal lines is smaller than the case when only one interconnect

20   switches ($T_a < T_b$).

The worst-case scenario for total delay and energy dissipation occurs when immediately adjacent lines simultaneously switch in opposite directions (time (c)), e.g. $\Delta V_1 = -V_{DD}$ and $\Delta V_2 = V_{DD}$. This leads to twice the delay increase due to capacitive

9

coupling compared to case b, as $|\Delta Q| = C_c \times |\Delta V_1 - \Delta V_2| = 2\,C_c\,V_{DD}$, and hence $T_c > T_b$. This worst case bus delay determines the bus cycle time and bus performance. Moreover, the energy dissipation due to capacitive coupling quadruples compared to case b, as $E_{Cc} = 0.5 \times C_c \times (\Delta V_1 - \Delta V_2)^2 = 2C_c V_{DD}^2 = 4E_{Cb}$.

5   The effect of the coupling capacitance changes according to the relative switching activities of the lines that it couples. The relative voltage difference $|\Delta V_1 - \Delta V_2|$ across the coupling capacitance C_c is 0, V_{DD} and $2V_{DD}$, in cases (a), (b), and (c), respectively. This indicates that the coupling capacitance had no effect in case (a) and twice the effect in case (c) compared to case (b). Thus, we can define an

10   equivalent coupling capacitance $C_{c,eq} = MCF \times C_c$, where MCF is the Miller coupling factor. MCF=0 when adjacent lines simultaneously switch similarly, MCF=1 when only one of the adjacent lines switch, and MCF=2 when adjacent lines simultaneously switch in opposite directions.

15   **Reducing the Average Energy Dissipation of the High-Performance Relative Delay Interconnect Architecture**

        One technique to reduce the worst-case delay is to introduce a relative delay between adjacent signal lines. This separates the simultaneous opposite switching edges (MCF=2) and transfers the switching to two consecutive cases of MCF=1.

20   This reduces the worst-case MCF from 2 to 1, hence reducing the worst-case total capacitance and thereby reducing the total bus delay.

        However, in arrangements introducing a relative delay to minimize the maximum MCF, the average energy dissipation remains the same. The reason is

that introducing a long enough relative delay between oppositely switching signal lines separates the simultaneous opposite switching edges (MCF=2) and transfers the switching to two consecutive cases of MCF=1 and thus reduces the energy dissipation of both signal lines due to coupling from $2C_cV_{DD}^2$ to

5     $2\times0.5C_cV_{DD}^2=C_cV_{DD}^2$. But on the other hand, introducing the same relative delay between similarly switching signal lines may increase the MCF from 0 to 1 and thus increase the energy dissipation of both signal lines due to coupling from 0 to $2\times0.5C_cV_{DD}^2=C_cV_{DD}^2$. Introducing a relative delay when MCF=1 does not change the MCF or the energy dissipation of the switching line. Thus, if a relative delay is

10   permanently introduced between any two adjacent signal lines, then the average energy dissipation may not change. It should also be noted that the energy dissipation penalty caused due to introducing the relative delay may occur based on the circuitry (i.e., additional buffers or skewed repeaters) performing the delay.

Embodiments of the present invention introduce a relative delay between

15   adjacent interconnects that not only reduces the maximum MCF and the worst-case delay, but also reduces the average bus energy dissipation. More specifically, embodiments of the present invention introduce a relative delay by lowering a supply voltage (hereafter V_{DD}) on alternate signal lines to $(V_{DD}-\Delta V)$. Lowering of the supply voltage on alternate signal lines increases the propagation delay of drivers and

20   repeaters on those lines, which lead to a relative delay between switching edges of neighboring signal lines and thus reducing the MCF of the bus/interconnect architecture. The average energy dissipation is proportional to the supply voltage. As discussed above, introducing a relative delay does not change the average

11

energy dissipation; however, embodiments of the present invention not only introduce a relative delay but also lower the supply voltage of alternate lines, and thus the average energy dissipation is reduced by $[1-(V_{DD}-\Delta V)^2/V_{DD}^2]$. Thus, the overall average energy dissipation of the bus is reduced by $0.5 \times [1-(V_{DD}-\Delta V)^2/V_{DD}^2]$.

5   FIG. 3 shows drivers for two adjacent signal lines of an interconnect architecture according to an example embodiment of the present invention. Other embodiments and configurations are also within the scope of the present invention. More specifically, FIG. 3 shows a first driver 400 and a second driver 420. As one example, the drivers 400 and 420 may correspond to the drivers 210 and 220, respectively, shown in FIG. 2A. Other drivers would be provided for additional signal

10   lines. The first driver 400 has a first transistor 402 and a second transistor 404 coupled between supply voltage V_{DD} and GROUND to provide a signal on a signal line 410 (or interconnect). The second driver 420 has a first transistor 422 and a second transistor 424 coupled between the supply voltage $V_{DD}-\Delta V$ and GROUND to

15   provide a signal on a signal line 430 (or interconnect). In this example, the signal lines 410 and 430 (such as corresponding to the interconnects 240 and 250, respectively in FIG. 2A) may extend in a parallel structure such that the signal line 430 is considered adjacent to the signal line 410. Other signal lines are not shown for ease of illustration.

20   The supply voltage of the second driver 420 is less than the supply voltage of the drivers 400 associated with the adjacent signal line 410. The lowering of the supply voltage on alternative lines (such as the supply voltage of the driver 420)

12

introduces a relative delay between adjacent signal lines, and thus reduces the average energy dissipation of the overall interconnect structure.

Combining Low Power Encoding and Dynamic Relative Delaying in Multi-Stage Time-Borrowing Buses

Both worst-case delay and average energy dissipation due to relative-switching of adjacent interconnects may be reduced in a different relative delay scheme. In this scheme, a dynamic delay may be introduced to various signal lines such that oppositely switching signal lines are relatively delayed while similarly switching signal lines are not relatively delayed. Thus, this scheme reduces the worst-case delay (MCF=2) to the delay of an MCF=1 scenario. In addition to that, this scheme reduces the average energy dissipation due to capacitive coupling in the case of MCF=2 from $2C_CV_{DD}^2$ to $C_CV_{DD}^2$ without increasing the energy dissipation due to capacitive coupling in the case of MCF=0. Thus, the overall average energy dissipation of the bus is significantly reduced by dynamically introducing relative delays between adjacent interconnects. A delay decision circuit may be provided that monitors an input data switching pattern and dynamically decides which interconnects should be delayed.

Low-power encoding techniques such as Bus-Invert coding may reduce the switching activity of individual lines. This reduces the energy dissipation associated with the self-switching of the bus lines. In order to reduce both the energy dissipation due to self-switching and that due to relative switching, low power coding techniques can be used followed by dynamic delaying. Applying a low-power bus coding

technique such as bus invert coding (by use of a bus invert encoder circuit) followed by dynamic delaying of signal lines (by a dynamic delay circuit) will lead to a significant energy dissipation reduction, but may add a delay due to the introduced delay as well as the delay associated with the bus invert coding hardware and the

5   delay decision circuit. However, when this architecture is applied to long multi-stage interconnect architectures with time-borrowing flip-flop (TBFF) circuits, then the delay overhead will be absorbed and amortized over all the stages of the interconnect architecture. Accordingly, this will achieve power reduction and significantly reduce the delay overhead per stage of the interconnect architecture.

10   FIG. 4 is a multi-stage interconnect architecture 500 according to an example embodiment of the present invention. Other embodiments and configurations are also within the scope of the present invention. More specifically, FIG. 4 shows a non-time borrowing flip-flop (NTBFF) circuit (or circuits) 510 coupled to a low power bus encoder such as the bus invert encoder circuit 520. The bus invert encoder

15   circuit 520 receives input data from the NTBFF circuit 510 and applies processing so as to minimize the amount of output switching across signal lines. The bus invert encoder circuit 520 is coupled to a dynamic delay circuit 530 that introduces a dynamic delay on certain signal lines based on an input data pattern. The dynamic delay circuit 530 outputs signals onto the signal lines. For ease of illustration, only

20   one signal line is shown in the figure for the first stage output from the dynamic delay circuit 530, although other signal lines would also be similarly coupled to the dynamic delay circuit 530 to form the first stage of the multi-stage interconnect architecture 500.

The multi-stage interconnect architecture 500 includes a first interconnect 540 and a second interconnect 550. Only two stages of the multi-stage interconnect are shown in FIG. 4 although numerous other stages may be provided in the area 555 for the multi-stage interconnect architecture 500. The multi-stage interconnect

5   architecture 500 includes a time-borrowing flip-flop (TBFF) circuit (or circuits) 560 between the various stages, such as between the first and second stages. The TBFF circuit 560 will allow portions of the delay to be passed to its subsequent bus stages. At the end of the multiple stages, such as after the area 555, the multi-stage interconnect architecture 500 also includes a non-time borrowing flip-flop (NTBFF)

10   circuit (or circuits) 570 and a low power bus decoder such as the bus invert decoder circuit 580. The bus invert decoder circuit 580 applies processing to decode the pattern of the bus invert encoder circuit 520 such that the proper output data pattern is provided.

When the above-described architecture is provided to long multi-stage

15   interconnects (or buses) with time-borrowing flip-flops circuits, then the delay overhead is absorbed and amortized (or averaged over all the stages of the interconnect). Thus, the above-described architecture provides more power dissipation reduction compared to only using dynamic relative delays, in addition to significantly reducing the delay overhead per stage.

20   Embodiments of the present invention provide low power bus coding followed by dynamic delaying the data of certain signal lines of a multi-stage interconnect structure. This leads to significant energy savings by reducing the energy dissipation due to switching while adding a small delay overhead per bus stage.

Embodiments of the present invention may be provided within one or more chips or chipsets within an electrical assembly. The term electrical assembly may be part of an electronic system. An electronic system may be any product including an electronic assembly. Examples of electronic systems include computers (e.g.

5   desktop, laptop, hand-held, server, etc.), wireless communications devices (e.g. cellular phone, cordless phones, pagers, etc.), computer-related peripherals (e.g., printers, scanners, monitors, etc.), entertainment devices (e.g. televisions, radios, stereos, tap and compact disc players, video cassette recorders, MP3 (Motion Picture Experts Group, Audio Layer 3) players, etc.), and the like.

10   Any reference in this specification to "one embodiment," "an embodiment," "example embodiment," etc., means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of such phrases in various places in the specification are not necessarily all referring to the same embodiment.

15   Further, when a particular feature, structure, or characteristic is described in connection with any embodiment, it is submitted that it is within the purview of one skilled in the art to affect such feature, structure, or characteristic in connection with other ones of the embodiments.

Although embodiments of the present invention have been described with
20   reference to a number of illustrative embodiments thereof, it should be understood that numerous other modifications and embodiments can be devised by those skilled in the art that will fall within the spirit and scope of the principles of this invention. More particularly, reasonable variations and modifications are possible in the

component parts and/or arrangements of the subject combination arrangement within the scope of the foregoing disclosure, the drawings and the appended claims without departing from the spirit of the invention. In addition to variations and modifications in the component parts and/or arrangements, alternative uses will also

5   be apparent to those skilled in the art.